Second- and Higher-Order Asymptotics For Erasure and List Decoding

Vincent Y. F. Tan and Pierre Moulin

arXiv:1402.4881v1 [cs.IT] 20 Feb 2014

Abstract

We derive the optimum second-order coding rates, known as second-order capacities, for erasure and list decoding. For erasure decoding for discrete memoryless channels, we show that second-order capacity is $\sqrt{V}\Phi^{-1}(\epsilon_{\mathrm{t}})$ where V is the channel dispersion and ϵ_{t} is the total error probability, i.e., the sum of the erasure and undetected errors. We show numerically that the expected rate at finite blocklength for erasures decoding can exceed the finite blocklength channel coding rate. We also show that the analogous result also holds for lossless source coding with decoder side information, i.e., Slepian-Wolf coding. For list decoding, we consider list codes of deterministic size that scales as $\exp(\sqrt{n}\,l)$ and show that the second-order capacity is $l+\sqrt{V}\Phi^{-1}(\epsilon)$ where ϵ is the permissible error probability. We also consider lists of polynomial size n^{α} and derive bounds on the third-order coding rate in terms of the order of the polynomial α. These bounds are tight for symmetric and singular channels. The direct parts of the coding theorems leverage on the simple threshold decoder and converses are proved using variants of the hypothesis testing converse.

I. Introduction

In many communication scenarios, it is advantageous to allow the decoder to have the option of either not deciding at all or putting out more than one estimate of the message. These are respectively known as *erasure* and *list* decoding respectively and have been studied extensively in the information theory literature [1]–[7]. The erasure and list options generally allow for smaller undetected error probabilities so these options are useful in practice.

In this paper, we revisit the problem or erasure and list decoding from the viewpoint of second- and third-order asymptotics. The study of second-order asymptotics for fixed (non-vanishing) error probability was first done by Strassen [8, Thm. 1.2] who showed for a well-behaved discrete memoryless channel (DMCs) W that the maximum number of codewords at (average or maximum) error probability ϵ, namely $M^*(W^n,\epsilon)$, satisfies

$$\log M^*(W^n,\epsilon) = nC + \sqrt{nV}\Phi^{-1}(\epsilon) + O(\log n), \tag{1}$$

where C and V are respectively the capacity and the dispersion of W and $\Phi^{-1}(\cdot)$ is the inverse of the standard Gaussian cumulative distribution function (cdf). Also see Kemperman [9, Thm. 11.3] for the corresponding result for constant composition codes. This line of work has been revisited by numerous authors recently and they considered various other channel models such as the additive white Gaussian (AWGN) channel [10], [11], bounds on the third-order (logarithmic) term [12]–[15] and extensions to multi-terminal problems such as the two-encoder Slepian-Wolf problem and the multiple-access channel [16].

A. Main Contributions

In this paper, for erasure decoding, we consider constant undetected and total (sum of undetected and erasure) error probabilities (numbers between 0 and 1) and we obtain the analogue of the second-order $\sqrt{n}$ term in (1). We show that the coefficient of the second-order term, termed the *second-order capacity*, is $\sqrt{V}\Phi^{-1}(\epsilon_{\mathrm{t}})$ where ϵ_{t} is the total error probability. The second-order capacity is thus completely independent of the undetected error probability. We then compute the expected rate at finite blocklength allowing erasures and show that it can exceed the finite blocklength rate without the erasure option. We show that these results carry over in a straightforward manner to the problem of lossless source coding with side information, i.e., the Slepian-Wolf problem [17] which was previously studied in the for the case without the erasure option [16].

V. .Y. F. Tan is with the Department of Electrical and Computer Engineering (ECE), National University of Singapore (Email: vtan@nus.edu.sg). P. Moulin is with the Department of Electrical and Computer Engineering (ECE), University of Illinois at Urbana-Champaign (Email: moulin@ifp.uiuc.edu).

For list decoding, we consider lists of deterministic size of order $\exp(\sqrt{n}\,l)$ and show that the second-order capacity is $l+\sqrt{V}\Phi^{-1}(\epsilon)$, where ϵ is the permissible error probability. We also consider lists of polynomial size n^{α} and demonstrate bounds on the third-order term in terms of the degree of the polynomial α. These bounds turn out to tight for channels that are symmetric and singular–a canonical example being the binary erasure channel [14]. To the best of the authors' knowledge, this is the first time that lists of size other than constant or exponential have been considered in the literature. Practically, the advantage of smaller lists is that they result in lower search complexity for the true message within the decoded list.

B. Related Work

Previously, the study of erasure and list decoding has been primarily from the error exponents perspective. We summarize some existing works here. Forney [1] derived optimal decision rules by generalizing the Neyman-Pearson lemma and also proved exponential upper bounds for the error probabilities using Gallager's Chernoff bounding techniques [18]. Shannon-Gallager-Berlekamp [2] proved exponential lower bounds for the error probabilities and also considered lists of exponential size $\exp(nl)$. They showed that sphere packing error exponent (evaluated at the code rate minus l) is an upper bound on the reliability function [19, Ex. 10.28]. Bounds for the error probabilities were derived by Telatar [3] using a general decoder parametrized by an asymmetric relation $\prec$ which is a function of the channel law. Blinovsky [4] studied the exponents of the list decoding problem at low (and even 0) rate. Csiszár-Körner [19, Thm. 10.11] present exponential upper bounds for universally attainable erasure decoding using the method of types. Moulin [5] generalized the treatment there and presented improved error exponents under some conditions. Recently, Merhav also considered alternative methods of analysis [6] and expurgated exponents for these problems [7]. The same author also derived erasure and list exponents for the Slepian-Wolf problem [20].

Another related line of work concerns constant error probability non-asymptotic fundamental limits of channel coding with various forms of feedback. Polyanskiy-Poor-Verdú [21] studied various *incremental redundancy* schemes and derived performance bounds under receiver- and transmitter-confirmation. In incremental redundancy systems/schemes, one is allowed to transmit a sequence of coded symbols with numerous opportunities for confirmation before the codeword is resent if necessary. In contrast in our study of channel coding with the erasure option in Sec. III-B, we analyze the expected performance of the easily-implementable Forney-style [1] *single-codeword repetition* scheme that allows for confirmation (or erasure) at the end of a complete codeword block and repeats the same message until the erasure event no longer occurs. We compare a quantity termed the *expected rate* and to the ordinary channel coding rate. Inspired by [21], Chen *et al.* [22] derived non-asymptotic bounds for feedback systems using incremental redundancy with noiseless transmitter confirmation. Williamson-Chen-Wesel [23] also improved on the bounds in [21] for variable-length feedback coding.

C. Structure of Paper

This paper is structured as follows: In Sec. II, we set the stage by introducing our notation and defining relevant quantities such as the second-order capacity. In Sec. III, we state our main results for channel coding with the decoding and list option. In Sec. IV, we show that the channel coding results carry over to the Slepian-Wolf problem [17]. All the proofs are detailed in Sec. V.

II. Problem Setting and Main Definitions

Let W be a random transformation (channel) from a discrete input alphabet $\mathcal{X}$ to a discrete output alphabet $\mathcal{Y}$. We denote length-n deterministic (resp. random) strings $\mathbf{x}=(x_1,\ldots,x_n)\in\mathcal{X}^n$ (resp. $\mathbf{X}=(X_1,\ldots,X_n)\in\mathcal{X}^n$) by lower case (resp. upper case) boldface. If W^n satisfies $W^n(\mathbf{y}|\mathbf{x})=\prod_{j=1}^n W(y_j|x_j)$ for every $(\mathbf{x},\mathbf{y})\in\mathcal{X}^n\times\mathcal{Y}^n$ and the sets $\mathcal{X}$ and $\mathcal{Y}$ are finite, W^n is said to be a DMC. We focus on DMCs in this paper but extensions to other channels such as the AWGN channel are straightforward. For a sequence $\mathbf{x}$, its *type* is the empirical distribution $P_{\mathbf{x}}(x)=\frac{1}{n}\sum_{i=1}^n \mathbf{1}\{x_i=x\}$. For a finite alphabet $\mathcal{X}$, let $\mathcal{P}(\mathcal{X})$ and $\mathcal{P}_n(\mathcal{X})$ be the set of probability mass functions and n-types [19] (types with denominator at most n) respectively. For two sequences $(\mathbf{x},\mathbf{y})\in\mathcal{X}^n\times\mathcal{Y}^n$, we say that U is a *conditional type* of $\mathbf{y}$ given $\mathbf{x}$ if $P_{\mathbf{x}}(x)U(y|x)=P_{\mathbf{x},\mathbf{y}}(x,y)$ for all $(x,y)\in\mathcal{X}\times\mathcal{Y}$. ($U$ is not unique if $P_{\mathbf{x}}(x)=0$ for some $x\in\mathcal{X}$.) The *U-shell* of a sequence $\mathbf{x}\in\mathcal{X}^n$, denoted as $\mathcal{T}_U(\mathbf{x})\subset\mathcal{Y}^n$ is the set of all $\mathbf{y}$ such that the joint type of $(\mathbf{x},\mathbf{y})$ is $P_{\mathbf{x}}\times U$. The set of all conditional types U for which $\mathcal{T}_U(\mathbf{x})$ is non-empty for some $\mathbf{x}$ with type P is denoted as $\mathcal{U}_n(\mathcal{Y};P)$. All logs are to the base 2 with the understanding that $\exp(t)=2^t$.

For information-theoretic quantities, we will mostly follow the notation in Csiszár and Körner [19]. We denote the *information capacity* of the DMC $W : \mathcal{X} \to \mathcal{Y}$ as $C = C(W) := \max_{P \in \mathcal{P}(\mathcal{X})} I(P, W)$. We let $\Pi = \Pi(W) := \{P \in \mathcal{P}(\mathcal{X}) : I(P, W) = C\}$ be the set of *capacity-achieving input distributions*. If (X, Y) has joint distribution $P(x)W(y|x)$, define $PW(y) := \sum_x P(x)W(y|x)$ to be the induced output distribution and

$$V(P, W) := \mathrm{E}_X \left[\mathsf{Var} \left(\log \frac{W(Y|X)}{PW(Y)} \,\middle|\, X \right) \right] \tag{2}$$

to be the *conditional information variance*. The ϵ-*dispersion* of the DMC W [8], [10], [11] is defined as

$$V_\epsilon := \begin{cases} V_{\min} := \min_{P \in \Pi} V(P, W) & \epsilon < 1/2 \\ V_{\max} := \max_{P \in \Pi} V(P, W) & \epsilon \geq 1/2 \end{cases}. \tag{3}$$

We will assume throughout that the DMC W satisfies $V_\epsilon > 0$. For integers $l \leq m$, we denote $[l : m] := \{l, l+1, \ldots, m\}$ and $[m] := [1 : m]$. Let $\Phi(x) := \int_{-\infty}^{x} (2\pi)^{-1/2} \mathrm{e}^{-t^2/2} \, \mathrm{d}t$ be the cdf of a standard Gaussian and $\Phi^{-1}(\cdot)$ its inverse. We now define erasure codes.

Definition 1. *An* M-erasure code *for* $W : \mathcal{X} \to \mathcal{Y}$ *is a pair of mappings* (f, φ) *such that* $f : [M] \to \mathcal{X}$ *and* $\varphi : \mathcal{Y} \to [0 : M]$. *The disjoint* decoding regions *are denoted as* $\mathcal{D}_m := \varphi^{-1}(m)$*; the* erasure region *is denoted as* $\mathcal{D}_0 := \varphi^{-1}(0)$*; and the* conditional undetected, erasure *and* total error probabilities *are defined as*

$$\lambda_{\mathrm{u}}(m) := \sum_{\tilde{m} \in [M] \setminus \{m\}} W(\mathcal{D}_{\tilde{m}} | f(m)) \tag{4}$$

$$\lambda_{\mathrm{e}}(m) := W(\mathcal{D}_0 | f(m)) \tag{5}$$

$$\lambda_{\mathrm{t}}(m) := \sum_{\tilde{m} \in [0:M] \setminus \{m\}} W(\mathcal{D}_{\tilde{m}} | f(m)) \tag{6}$$

Note that $\lambda_{\mathrm{u}}(m) + \lambda_{\mathrm{e}}(m) = \lambda_{\mathrm{t}}(m)$. Typically, the code is designed so that $\lambda_{\mathrm{u}}(m) \ll \lambda_{\mathrm{e}}(m)$ as the cost of making an undetected error is much higher than that of declaring an erasure.

Definition 2. *An* $(M, \epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}})_{\mathrm{a,b}}$-erasure code *for* W *is an* M*-erasure code for the same channel where*

1) *If* $(\mathrm{a}, \mathrm{b}) = (\max, \max)$,
$$\max_{m \in [M]} \lambda_{\mathrm{u}}(m) \leq \epsilon_{\mathrm{u}}, \quad \max_{m \in [M]} \lambda_{\mathrm{t}}(m) \leq \epsilon_{\mathrm{t}}. \tag{7}$$

2) *If* $(\mathrm{a}, \mathrm{b}) = (\max, \mathrm{ave})$
$$\max_{m \in [M]} \lambda_{\mathrm{u}}(m) \leq \epsilon_{\mathrm{u}}, \quad \frac{1}{M} \sum_{m \in [M]} \lambda_{\mathrm{t}}(m) \leq \epsilon_{\mathrm{t}}. \tag{8}$$

3) *If* $(\mathrm{a}, \mathrm{b}) = (\mathrm{ave}, \max)$,
$$\frac{1}{M} \sum_{m \in [M]} \lambda_{\mathrm{u}}(m) \leq \epsilon_{\mathrm{u}}, \quad \max_{m \in [M]} \lambda_{\mathrm{t}}(m) \leq \epsilon_{\mathrm{t}}. \tag{9}$$

4) *If* $(\mathrm{a}, \mathrm{b}) = (\mathrm{ave}, \mathrm{ave})$,
$$\frac{1}{M} \sum_{m \in [M]} \lambda_{\mathrm{u}}(m) \leq \epsilon_{\mathrm{u}}, \quad \frac{1}{M} \sum_{m \in [M]} \lambda_{\mathrm{t}}(m) \leq \epsilon_{\mathrm{t}}. \tag{10}$$

In Definition 2, we consider erasure codes with constraints on the *undetected* and *total* error probabilities similar to [1], [5], [6]. An alternate formulation would be to consider $(M, \epsilon_{\mathrm{u}}, \epsilon_{\mathrm{e}})_{\mathrm{a,b}}$-erasure codes where ϵ_{e} is the *erasure* probability. We find the former formulation more traditional and the analysis is also somewhat easier.

Definition 3. *A number* $r \in \mathbb{R}$ *is an* $(\epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}})_{\mathrm{a,b}}$-achievable erasure second-order coding rate *for the DMC* W^n *with capacity* C *if there exists a sequence of* $(M_n, \epsilon_{\mathrm{u},n}, \epsilon_{\mathrm{t},n})_{\mathrm{a,b}}$-erasure codes *such that*

$$\liminf_{n \to \infty} \frac{1}{\sqrt{n}} (\log M_n - nC) \geq r, \tag{11}$$

$$\limsup_{n \to \infty} \epsilon_{\mathrm{u},n} \leq \epsilon_{\mathrm{u}}, \quad \textit{and} \tag{12}$$

$$\limsup_{n \to \infty} \epsilon_{\mathrm{t},n} \leq \epsilon_{\mathrm{t}}. \tag{13}$$

The $(\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})_{\mathrm{a,b}}$-erasure second-order capacity $r^*_{\mathrm{era,a,b}}(\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})$ *is the supremum of all* $(\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})_{\mathrm{a,b}}$*-achievable erasure second-order coding rates.*

We now turn our attention to codes which allow their decoders to output a *list* of messages. Let $\binom{[M]}{j}$ be the set of subsets of $[M]$ of size j. Furthermore, we use the notation $\binom{[M]}{\leq L} := \cup_{0\leq j\leq L}\binom{[M]}{j}$ to denote the set of subsets of $[M]$ of size not exceeding L.

Definition 4. *An* (M,L)-list code *for* $W:\mathcal{X}\to\mathcal{Y}$ *is a pair of mappings* (f,φ) *such that* $f:[M]\to\mathcal{X}$ *and* $\varphi:\mathcal{Y}\to\binom{[M]}{\leq L}$. *The (not-necessarily disjoint)* decoding regions *are denoted as* $\mathcal{D}_m := \{y\in\mathcal{Y}: m\in\varphi(y)\}$ *and the* conditional error probability *is defined as*

$$\lambda(m) := W(\mathcal{Y}\setminus\mathcal{D}_m|f(m)). \tag{14}$$

Definition 5. *An* $(M,L,\epsilon)_{\mathrm{a}}$-list code *for* W *is an* (M,L)*-list code for the same channel where if* $\mathrm{a}=\max$*,*

$$\max_{m\in[M]}\lambda(m)\leq\epsilon, \tag{15}$$

or if $\mathrm{a}=\mathrm{ave}$*,*

$$\frac{1}{M}\sum_{m\in[M]}\lambda(m)\leq\epsilon. \tag{16}$$

Definition 6. *A number* $r\in\mathbb{R}$ *is an* $(l,\epsilon)_{\mathrm{a}}$-achievable list second-order coding rate *for the DMC* W^n *with capacity* C *if there exists a sequence of* $(M_n,L_n,\epsilon_n)_{\mathrm{a}}$-list codes *such that in addition to* (11)*, the following hold*

$$\limsup_{n\to\infty}\frac{1}{\sqrt{n}}\log L_n\leq l,\quad \textit{and} \tag{17}$$

$$\limsup_{n\to\infty}\epsilon_n\leq\epsilon. \tag{18}$$

The $(l,\epsilon)_{\mathrm{a}}$-list second-order capacity $r^*_{\mathrm{list,a}}(l,\epsilon)$ *is the supremum of all* $(l,\epsilon)_{\mathrm{a}}$*-achievable list second-order coding rates.*

According to (17), we stipulate that the list size grows as $\exp(\sqrt{n}\,l)$. This differs from previous works on list decoding in which the list size is either constant [7], [18] or exponential [2]–[7], [19], [20]. This scaling affects the second-order (dispersion) term. To understand how the list size may affect the higher-order term, consider the following definition.

Definition 7. *A number* $s\in\mathbb{R}$ *is an* (α,ϵ)-achievable list third-order coding rate *for the DMC* W^n *with capacity* C *and positive* ϵ*-dispersion* V_ϵ *if there exists a sequence of* $(M_n,L_n,\epsilon_n)_{\mathrm{ave}}$-list codes *such that*

$$\liminf_{n\to\infty}\frac{1}{\log n}\big(\log M_n - nC - \sqrt{nV_{\epsilon_n}}\Phi^{-1}(\epsilon_n)\big)\geq s, \tag{19}$$

$$\limsup_{n\to\infty}\frac{\log L_n}{\log n}\leq\alpha,\quad \textit{and} \tag{20}$$

$$\limsup_{n\to\infty}\sqrt{n}(\epsilon_n-\epsilon)<\infty. \tag{21}$$

The (α,ϵ)-list third-order capacity $s^*_{\mathrm{list}}(\alpha,\epsilon)$ *is the supremum of all* (α,ϵ)*-achievable list third-order coding rates.*

Inequality (20) implies that the size of the list grows polynomially and in particular it scales as $O(n^\alpha)$. This scaling affects the third-order (logarithmic) term studied by a number of authors [12]–[15] in the context of ordinary channel coding (without list decoding). By (19), if $s\in\mathbb{R}$ is (α,ϵ)-achievable, there exists a sequence of codes of sizes M_n satisfying

$$\log M_n\geq nC+\sqrt{nV_{\epsilon_n}}\Phi^{-1}(\epsilon_n)+s\log n+o(\log n), \tag{22}$$

and having list sizes $L_n=O(n^\alpha)$ and average error probabilities ϵ_n not exceeding $\epsilon+O(n^{-1/2})$. The more stringent condition on the sequence of error probabilities $\{\epsilon_n\}_{n\geq1}$ in (21) (relative to (18)) is because ϵ_n appears as the argument in the $\Phi^{-1}(\cdot)$ function, which forms part of the coefficient of the $\sqrt{n}$ second-order term in the asymptotic expansion of $\log M_n$. If the weaker condition (18) were in place, the approximation error between ϵ_n

and the target ϵ which is of the order $o(\sqrt{n})$ would affect the third-order term, which is the object of study here. The stronger condition in (21) ensures that the third-order (logarithmic) term is unaffected by the approximation error between ϵ_n and ϵ which is now of the order $O(1)$.

III. Main Results for Channel Coding

In this section, we summarize the main results of this paper concerning channel coding with an erasure or list option. For simplicity, we assume that the DMC W satisfies $V_\epsilon > 0$ though our results can be extended in a straightforward manner to the case where $V_\epsilon = 0$.

A. Decoding with Erasure Option

Theorem 1. *For any* $0 \leq \epsilon_{\mathrm{u}} < \epsilon_{\mathrm{t}} < 1$,

$$r^*_{\mathrm{era,a,b}}(\epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}}) = \sqrt{V_{\epsilon_{\mathrm{t}}}}\Phi^{-1}(\epsilon_{\mathrm{t}}), \tag{23}$$

where (a, b) *can be any element in* $\{\max, \mathrm{ave}\}^2$.

The proof of Theorem 1 can be found in Section V-A.

A few comments are in order: First, Theorem 1 implies that if $M^*(W^n; \epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}})$ is the maximum number of codewords that can be transmitted over W^n with undetected and total error ϵ_{u} and ϵ_{t} respectively, then

$$\log M^*(W^n; \epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}}) = nC + \sqrt{nV_{\epsilon_{\mathrm{t}}}}\Phi^{-1}(\epsilon_{\mathrm{t}}) + o(\sqrt{n}). \tag{24}$$

We see that the backoff from the capacity at blocklength n is approximately $-\sqrt{V_{\epsilon_{\mathrm{t}}}/n}\,\Phi^{-1}(\epsilon_{\mathrm{t}})$ independent of ϵ_{u}. (This backoff is positive for $\epsilon_{\mathrm{t}} < 1/2$.) Observe that the second-order term does not depend on ϵ_{u}, the undetected error probability. Only the total error probability comes into play in the asymptotic characterization of $\log M^*(W^n; \epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}})$. In fact, in the proof, we first argue that it suffices to show that $\sqrt{V_{\epsilon_{\mathrm{e}}}}\Phi^{-1}(\epsilon_{\mathrm{e}})$ is an achievable $(0, \epsilon_{\mathrm{e}})_{\mathrm{ave,ave}}$-erasure second-order coding rate for W^n, i.e., the undetected error probability is asymptotically 0. Clearly, any achievable $(0, \epsilon_{\mathrm{t}})_{\mathrm{ave,ave}}$-erasure second-order coding rate is also an achievable $(\epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}})_{\mathrm{ave,ave}}$-erasure second-order coding rate for any $\epsilon_{\mathrm{u}} \in [0, \epsilon_{\mathrm{t}})$.

Second, in the direct part of the proof of Theorem 1, we use threshold decoding, i.e. declare that message m is sent if the empirical mutual information is higher than a threshold. If no message's empirical mutual information exceeds the threshold, then an erasure is declared. This simple rule, though *universal*, is not the optimal one (in terms of minimizing the total error while holding the undetected error fixed). The optimal rule was derived using a generalized version of the Neyman-Pearson lemma by Forney [1, Thm. 1] and it is stated as

$$\mathcal{D}^*_m := \left\{\mathbf{y} : W^n(\mathbf{y}|f(m)) \geq \psi \sum_{\tilde{m}\in[M]\setminus\{m\}} W^n(\mathbf{y}|f(\tilde{m}))\right\}, \tag{25}$$

for some threshold $\psi > 0$. However, this rule appears to be difficult for second-order analysis and is more amenable to error exponent analysis [5], [7]. Because the likelihood of the second most likely codeword is usually much higher than the rest (excluding the first), Forney also suggested the simpler but, in general, suboptimal rule [1, Eq. (11a)]

$$\mathcal{D}'_m := \left\{\mathbf{y} : W^n(\mathbf{y}|f(m)) \geq \psi \max_{\tilde{m}\in[M]\setminus\{m\}} W^n(\mathbf{y}|f(\tilde{m}))\right\}. \tag{26}$$

We analyzed this rule in the asymptotic setting (i.e., as n tends to infinity) in the same way as one analyzes the random coding union (RCU) bound [10, Thm. 16] [12, Sec. 7] for ordinary channel coding but the analysis is more involved than threshold decoding which suffices for proving the second-order achievability of (23).

Third, the converse is based on Strassen's idea [8, Eq. (4.18)], establishing a clear link between point-to-point channel coding and binary hypothesis testing. Also see Kemperman's general converse bound [9, Lem. 3.1] and for a more modern treatment, the various forms of the *meta-converse* in [10, Sec. III-E]. The hypothesis testing converse technique only depends on the total error probability, explaining the presence of ϵ_{t} and not ϵ_{u} in (23).

Finally, we remark that Theorem 1 (as well as Theorem 2 to follow bar the statement in (37)) carries over verbatim for the AWGN channel where the capacity and dispersion are $C = \frac{1}{2}\log(1+S)$ and $V = \log^2 \mathrm{e} \cdot \frac{S(S+2)}{2(S+1)^2}$ respectively and S is the signal-to-noise ratio (SNR) of the AWGN channel.

B. The Expected Rate and An Example

We now compare the expected rate achieved using decoding with the erasure option to ordinary channel coding. Define

$$\epsilon_{\mathrm{e}} := \epsilon_{\mathrm{t}} - \epsilon_{\mathrm{u}} \tag{27}$$

as the *erasure probability* and note from the assumption of Theorem 1 that $\epsilon_{\mathrm{e}} > 0$. Now consider sending $b \in \mathbb{N}$ *independent* blocks of information each of length $n \in \mathbb{N}$. Because transmission succeeds (no erasure declared) with probability $1 - \epsilon_{\mathrm{e}}$, the total number of bits we can transmit in each block is well approximated by the random variable

$$R_{\mathrm{e}}^{(n)} := \begin{cases} C + \sqrt{V_{\epsilon_{\mathrm{t}}}/n}\,\Phi^{-1}(\epsilon_{\mathrm{t}}) & \text{w.p. } 1-\epsilon_{\mathrm{e}} \\ 0 & \text{w.p. } \epsilon_{\mathrm{e}} \end{cases} \tag{28}$$

This random variable has expectation

$$\mathrm{E}\big[R_{\mathrm{e}}^{(n)}\big] := (1-\epsilon_{\mathrm{e}})\left(C + \sqrt{\frac{V_{\epsilon_{\mathrm{t}}}}{n}}\Phi^{-1}(\epsilon_{\mathrm{t}})\right). \tag{29}$$

Fix $\delta \in (0, \min\{\epsilon_{\mathrm{e}}, 1-\epsilon_{\mathrm{e}}\})$. By Hoeffding's inequality, the total number of bits we can transmit over the b blocks is in the interval

$$\left[(1-\epsilon_{\mathrm{e}}-\delta)\cdot b\cdot\left(nC + \sqrt{nV_{\epsilon_{\mathrm{t}}}}\Phi^{-1}(\epsilon_{\mathrm{t}})\right), (1-\epsilon_{\mathrm{e}}+\delta)\cdot b\cdot\left(nC + \sqrt{nV_{\epsilon_{\mathrm{t}}}}\Phi^{-1}(\epsilon_{\mathrm{t}})\right)\right] \tag{30}$$

with probability exceeding $1 - 2\,\mathrm{e}^{-b\delta^2}$. This reduction in rate in (29)–(30) by the factor of $1-\epsilon_{\mathrm{e}}$ in the so-called *single-codeword repetition* scheme was first observed by Forney [1, Eq. (49)]. Essentially, one may use *decision feedback* to resend the entire block of information if there is an erasure.

For ordinary channel coding with error probability ϵ_{u}, we can send approximately

$$b\cdot\left(nC + \sqrt{nV_{\epsilon_{\mathrm{u}}}}\Phi^{-1}(\epsilon_{\mathrm{u}})\right) \tag{31}$$

bits over the b independent blocks and so, dividing by nb, the non-asymptotic channel coding rate (analogue of (29)) can be approximated [10] by

$$R_{\mathrm{c}}^{(n)} := C + \sqrt{\frac{V_{\epsilon_{\mathrm{u}}}}{n}}\Phi^{-1}(\epsilon_{\mathrm{u}}). \tag{32}$$

In the analysis in (28)–(32), we have assumed that the Gaussian approximation is sufficiently accurate. It was numerically shown in [10] that the Gaussian approximation is accurate for some channels (such as the binary symmetric channel, binary erasure channel and additive white Gaussian noise channel) and moderate blocklengths (of order ≈ 100) and error probabilities (or order $\approx 10^{-6}$). Also see Fig. 1 for a precise quantification of the accuracy of the Gaussian approximation in our setting.

Clearly if $0 \le \epsilon_{\mathrm{u}} < \epsilon_{\mathrm{e}} \le \epsilon_{\mathrm{t}} < 1$ are constants,

$$C = \lim_{n\to\infty} R_{\mathrm{c}}^{(n)} > \lim_{n\to\infty} \mathrm{E}\big[R_{\mathrm{e}}^{(n)}\big] = (1-\epsilon_{\mathrm{e}})C, \tag{33}$$

so there is no advantage in allowing for erasures asymptotically. However, in finite blocklength (by this we mean the per-block blocklength n) regime, for "moderate" ϵ_{e}, we may have

$$R_{\mathrm{c}}^{(n)} < \mathrm{E}\big[R_{\mathrm{e}}^{(n)}\big] \tag{34}$$

so erasure decoding may be advantageous *in expectation*. We illustrate the difference between $\mathrm{E}\big[R_{\mathrm{e}}^{(n)}\big]$ and $R_{\mathrm{c}}^{(n)}$ with a concrete example.

Example: In Fig. 1, we consider a binary symmetric channel (BSC) with crossover probability $q = 0.11$ so $C = 0.5$ bits/channel use and $V = 0.891$ bits^2/channel use. (For the BSC, V_ϵ does not depend on ϵ.) We keep the undetected error probability at $\epsilon_{\mathrm{u}} = 10^{-6}$ and vary the erasure error probability ϵ_{e} in the interval $[10^{-6}, 10^{-1}]$. We chose two blocklengths $n \in \{2\times 10^3, 1\times 10^5\}$. We observe that for blocklength $n = 2\times 10^3$ and a moderate erasure probability of $\epsilon_{\mathrm{e}} \approx 10^{-2}$, the gain of coding with the erasure option over ordinary channel coding is rather pronounced. This can be seen by comparing either the Gaussian approximations or the finite blocklength bounds. This gain is reduced if (i) ϵ_{e} is increased because we retransmit the whole block more often on average via the use of decision feedback or (ii) n becomes large so the second-order $\sqrt{n}$ term becomes less significant (cf. (33)).

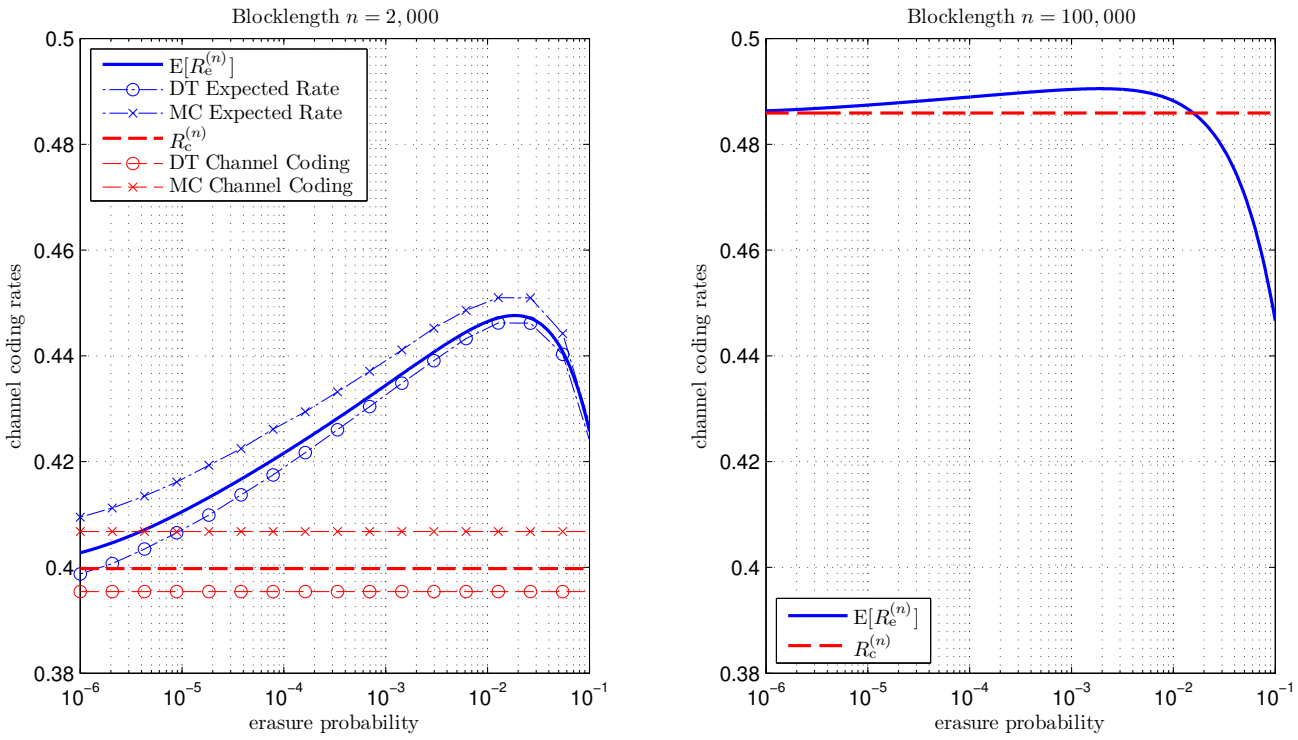

Fig. 1. Comparison of non-asymptotic rates and Gaussian approximationss for channel coding with and without the erasure option. Observe that $\mathrm{E}[R_{\mathrm{e}}^{(n)}]$ can be larger $R_{\mathrm{c}}^{(n)}$ for finite n. On the left plot, DT Expected Rate and MC Expected Rate respectively stand for the dependence-testing (DT) [10, Thm. 34] and meta-converse (MC) [10, Thm. 35] finite blocklength bounds for the expected rate with the given parameters $(\epsilon_{\mathrm{e}}, \epsilon_{\mathrm{u}}, n, q)$. DT Channel Coding and MC Channel Coding respectively stand for the DT and MC finite blocklength bounds for channel coding with ordinary decoding with the given parameters $(\epsilon_{\mathrm{u}}, n, q)$. The finite blocklength bounds are difficult to compute numerically for large blocklengths and small error probabilities and so they are not shown on the right plot.

We also note from the left plot that the Gaussian approximation is reasonably accurate when compared to the dependence-testing [10, Thm. 34] and meta-converse [10, Thm. 35] finite blocklength bounds under the current settings. The DT bound is especially close to the Gaussian approximation when the performance of coding with erasures peaks.

Finally, we remark that the advantage of coding with the erasures option over ordinary channel coding was also shown from the error exponents perspective by Forney in [1]. More precisely, Forney [1, Eq. (55)] showed that the feedback exponent has slope -1 for rates near and below capacity, thus improving on the ordinary exponent which has slope 0 in the same region.

C. List Decoding

Theorem 2. *For any $0 < \epsilon < 1$, we have the second-order result*

$$r^*_{\mathrm{list,a}}(l, \epsilon) = l + \sqrt{V_\epsilon}\Phi^{-1}(\epsilon). \tag{35}$$

where $\mathrm{a} \in \{\max, \mathrm{ave}\}$. Furthermore, if $\alpha \geq 0$ and $0 < \epsilon < 1$, we also have the third-order result

$$\alpha \leq s^*_{\mathrm{list}}(\alpha, \epsilon) \leq \frac{1}{2} + \alpha. \tag{36}$$

If the DMC W is symmetric in the Gallager sense[1] *and singular,*[2] *then we can make the stronger statement*

$$s^*_{\mathrm{list}}(\alpha, \epsilon) = \alpha. \tag{37}$$

[1]A DMC is *symmetric in the Gallager sense* [18, Sec. 4.5, pp. 94] if the set of channel outputs $\mathcal{Y}$ can be partitioned into subsets such that within each subset, the matrix of transition probabilities satisfies the following: every row (resp. column) is a permutation of every other row (resp. column).

[2]A DMC W is *singular* if for all $(x, y, \bar{x}) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{X}$, with $W(y|x)W(y|\bar{x}) > 0$ it is true that $W(y|x) = W(y|\bar{x})$.

The proof of Theorem 2 which is partly based on Proposition 3 below can be found in Section V-B.

Theorem 2 shows that if we allow the list to grow as $\exp(\sqrt{n}\,l)$, then the second-order capacity is increased by l. This is concurs with the intuition we obtain from the analysis of error exponents for list decoding by Shannon-Gallager-Berlekamp [2]. See also the exercises in Gallager [18, Ex. 5.20] and Csiszár-Körner [19, Ex. 10.28] which are concerned solely with first-order (capacity) analysis.

For the third-order result in (36), we observe that there is a gap. This is, in part, due to the fact that we use threshold decoding to decide which messages belong to the list. This decoding strategy appears to be suboptimal in the third-order sense for DMCs [12, Sec. 7] [24, Thm. 53] and AWGN channels [15]. It also appears that one needs to use a version of maximum-likelihood (ML) decoding as in (26) and analyze an analogue of the RCU bound [10, Thm. 16] carefully to obtain the additional $\frac{1}{2}$ for the direct part. Whether this can be done for channel coding with list decoding to obtain a tight third-order result general DMCs is an open question. Nonetheless for singular and symmetric channels considered by Altuğ and Wagner [14], such as the BEC, the converse (upper bound) can be tightened and this results in the conclusive result in (37).

A by-product of the proof of Theorem 2 is the following non-asymptotic converse bound for list-decoding which may be of independent interest.

Proposition 3. *Let $\beta_\alpha(P,Q)$ be the best (smallest) type-II error in a (deterministic) hypothesis test between P and Q subject to the type-I error being no larger than $1-\alpha$. Every $(M,L,\epsilon)_{\text{ave}}$-list code for $W:\mathcal{X}\to\mathcal{Y}$ satisfies*

$$\frac{M}{L} \le \inf_{Q\in\mathcal{P}(\mathcal{Y})} \sup_{P\in\mathcal{P}(\mathcal{X})} \frac{1}{\beta_{1-\epsilon}(P\times W, P\times Q)}. \tag{38}$$

This bound immediately reduces to the so-called *meta-converse* in [10, Sec. III-E] by setting $L=1$. We will see that the *ratio* M/L plays a critical role in both the converse and direct parts. This is also evident in existing works such as Shannon-Gallager-Berlekamp [2, Sec. IV] but the non-asymptotic bound in (38) appears to be novel.

IV. An Extension to Slepian-Wolf Coding

In this section, we show that the techniques developed are also applicable to lossless source coding with decoder side information, i.e., the Slepian-Wolf problem [17]. Let $P_{XY}\in\mathcal{P}(\mathcal{X}\times\mathcal{Y})$ be a correlated source where the alphabets $\mathcal{X}$ and $\mathcal{Y}$ are finite. The assumption that $\mathcal{X}$ and $\mathcal{Y}$ are finite sets can be dispensed at the cost of non-universality in the coding scheme [25].

Definition 8. *An* M-code *for the correlated source P_{XY} is defined by two functions $f:\mathcal{X}\to[M]$ and $\varphi:[M]\times\mathcal{Y}\to\mathcal{X}\cup\{\emptyset\}$ where $\emptyset$ represents the erasure symbol.*

Definition 9. *An* $(M,\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})$-code *for the correlated source P_{XY} is an M-code satisfying*

$$\sum_{x,y} P_{XY}(x,y)\mathbf{1}\big\{\varphi(f(x),y)\in\mathcal{X}\setminus\{x\}\big\}\le\epsilon_{\mathrm{u}}, \quad \textit{and} \tag{39}$$

$$\sum_{x,y} P_{XY}(x,y)\mathbf{1}\big\{\varphi(f(x),y)\neq x\big\}\le\epsilon_{\mathrm{t}}. \tag{40}$$

The parameters ϵ_{u} and ϵ_{t} are known as the undetected *and* total error probabilities *respectively.*

We consider discrete, stationary and memoryless sources $P_{\mathbf{XY}}(\mathbf{x},\mathbf{y})=\prod_{j=1}^n P_{XY}(x_j,y_j)$ in which the source alphabet is $\mathcal{X}^n$ and the side-information alphabet is $\mathcal{Y}^n$.

Definition 10. *A number $r\in\mathbb{R}$ is said to be an* $(\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})$-achievable second-order coding rate *for the correlated source $\{P_{X^nY^n}\}_{n\ge1}$ if there exists a sequence of $(M_n,\epsilon_{\mathrm{u},n},\epsilon_{\mathrm{t},n})$-codes such that*

$$\limsup_{n\to\infty}\frac{1}{\sqrt{n}}(\log M_n - nH(X|Y))\le r \tag{41}$$

$$\limsup_{n\to\infty}\epsilon_{\mathrm{u},n}\le\epsilon_{\mathrm{u}} \tag{42}$$

$$\limsup_{n\to\infty}\epsilon_{\mathrm{t},n}\le\epsilon_{\mathrm{t}}. \tag{43}$$

The infimum of all $(\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})$-achievable second-order coding rates is the optimum second-order coding rate $r^*_{\mathrm{era}}(\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})$.

Hence, we are allowing the erasure option for the Slepian-Wolf problem but we restrict the undetected and total errors to be at most ϵ_{u} and ϵ_{t} respectively. The second-order asymptotics for the Slepian-Wolf problem with two encoders without erasures was studied by Tan and Kosut in [16].

Define $V(X|Y) := \mathsf{Var}(-\log P_{X|Y}(X|Y))$ to be the *conditional varentropy* [26] of the stationary, memoryless source P_{XY}. The following result here parallels Theorem 1 pertaining to channel coding with the erasure option.

Theorem 4. *For any* $0 \leq \epsilon_{\mathrm{u}} < \epsilon_{\mathrm{t}} < 1$,

$$r^*_{\mathrm{era}}(\epsilon_{\mathrm{u}}, \epsilon_{\mathrm{t}}) = \sqrt{V(X|Y)}\Phi^{-1}(1-\epsilon_{\mathrm{t}}). \tag{44}$$

The proof of Theorem 4 can be found in Section V-C. The direct part uses random binning [27] and thresholding of the empirical conditional entropy [16] and the converse uses a non-asymptotic information spectrum converse bound by Miyake and Kanaya [28]. Also see [29, Thm. 7.2.2].

Again, we observe that the optimum second-order coding rate does not depend on the undetected error probability ϵ_{u} as long as it is smaller than the total error probability ϵ_{t}. Hence, the observation we made in Fig. 1–namely that at finite blocklengths the expected performance with the erasure option can exceed that without the erasure option–also applies in the Slepian-Wolf setting.

V. Proofs

In this section, we provide the proofs of the theorems in the paper.

A. Decoding with Erasure Option: Proof of Theorem 1

1) Converse Part: Recall that $\beta_\alpha(P,Q)$ is the best (smallest) type-II error in a hypothesis test (without randomization) of P versus Q subject to the condition that the type-I error is no larger than $1-\alpha$. This function was studied extensively in [30]. Let us fix any $(M_n, \epsilon_{\mathrm{u},n}, \epsilon_{\mathrm{t},n})_{\max,\max}$-erasure code for W^n. This means that

$$\min_{m\in[M_n]} W^n(\mathcal{D}_m|f(m)) \geq 1-\epsilon_{\mathrm{t},n}. \tag{45}$$

Note that only the total error comes into play in (45) and thus the second-order capacity in (23) only depends on ϵ_{t}. In addition, we assume that the code is constant composition, i.e. all codewords are of the same type P. This only leads to a $O(\log n)$ penalty in $\log M_n$ which does not affect the second-order term. Now, for any permutation invariant output distribution $Q \in \mathcal{P}(\mathcal{Y}^n)$ (this means that $Q(y_1,\ldots,y_n) = Q(y_{\pi(1)},\ldots,y_{\pi(n)})$ for every permutation $\pi:[n]\to[n]$), we have

$$\begin{aligned}
1 &\geq \sum_{m=1}^{M_n} Q(\mathcal{D}_m) && (46)\\
&\geq \sum_{m=1}^{M_n} \min_{\substack{\mathcal{D}\subset\mathcal{Y}^n:\\ W^n(\mathcal{D}|f(m))\geq 1-\epsilon_{\mathrm{t},n}}} Q(\mathcal{D}) && (47)\\
&= \sum_{m=1}^{M_n} \beta_{1-\epsilon_{\mathrm{t},n}}(W^n(\cdot|f(m)), Q) && (48)\\
&= M_n \beta_{1-\epsilon_{\mathrm{t},n}}(W^n(\cdot|\mathbf{x}), Q), && (49)
\end{aligned}$$

where (46) follows because $\mathcal{D}_m, m\in[M_n]$ are disjoint; (47) follows from (45); (48) uses the definition of β_α; and finally (49) follows from the fact that $\beta_{1-\epsilon_{\mathrm{t},n}}(W^n(\cdot|f(m)),Q)$ does not depend on m for permutation invariant Q (which is what we choose Q to be) and constant composition codes [8]. We also used $\mathbf{x}$ to denote any element in the type class $\mathcal{T}_P$. Choose Q to be the product distribution $(PW)^n$. Since $\epsilon_{\mathrm{t},n}$ satisfies (13), for every $\eta\in(0,1-\epsilon_{\mathrm{t}})$, there exists sufficiently large n such that $\epsilon_{\mathrm{t},n} \leq \epsilon_{\mathrm{t}}+\eta$. Hence, from (49),

$$M_n \leq \beta^{-1}_{1-\epsilon_{\mathrm{t}}-\eta}(W^n(\cdot|\mathbf{x}), Q). \tag{50}$$

Now, we use [8, Sec. 2] to assert that for all $\epsilon\in(0,1)$,

$$-\log\beta_{1-\epsilon}(W^n(\cdot|\mathbf{x}),Q) = nI(P,W) + \sqrt{nV(P,W)}\Phi^{-1}(\epsilon) + O(\log n). \tag{51}$$

Hence, by putting (50) and (51) together, we obtain

$$\log M_n \leq \max_{P\in\mathcal{P}(\mathcal{X})} nI(P,W) + \sqrt{nV(P,W)}\Phi^{-1}(\epsilon_{\mathrm{t}}+\eta) + O(\log n) \tag{52}$$

By using the usual continuity arguments (e.g. [13, Lem. 7]),

$$\log M_n \leq nC + \sqrt{nV_{\epsilon_{\mathrm{t}}+\eta}}\Phi^{-1}(\epsilon_{\mathrm{t}}+\eta) + O(\log n). \tag{53}$$

Thus, by letting $n\to\infty$,

$$\limsup_{n\to\infty} \frac{1}{\sqrt{n}}(\log M_n - nC) \leq \sqrt{V_{\epsilon_{\mathrm{t}}+\eta}}\Phi^{-1}(\epsilon_{\mathrm{t}}+\eta). \tag{54}$$

The converse proof is complete for the $(\max,\max)$ case by taking $\eta\downarrow 0$. Note that even though V_ϵ is, in general, discontinuous at $\epsilon=\frac{1}{2}$ (i.e., when $V_{\min}<V_{\max}$), we may let $\eta\downarrow 0$ and use the fact that $\Phi^{-1}(\frac{1}{2})=0$ to assert that $\lim_{\eta\to 0}\sqrt{V_{\frac{1}{2}+\eta}}\Phi^{-1}(\frac{1}{2}+\eta)=0$ as desired.

To get to the $(\mathrm{ave},\mathrm{ave})$ setting, first we let $M^*_{\mathrm{a,b}}(W,\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})$ be the maximum number of codewords M in an erasure with undetected and total error probabilities $\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}}$ respectively under the (a,b)-setting. By using an expurgation argument as in [10, Eq. (284)] it is easy to show that for all $\tau>1$,

$$M^*_{\mathrm{ave,ave}}(W^n,\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}}) \leq \left(\frac{1}{1-1/\tau}\right)^2 M^*_{\max,\max}(W^n,\tau\epsilon_{\mathrm{u}},\tau\epsilon_{\mathrm{t}}). \tag{55}$$

Setting $\tau := 1+\frac{1}{\sqrt{n}}$ proves the claim for all the other 3 cases.

2) Direct Part: It suffices to prove that $\sqrt{V_{\epsilon_{\mathrm{e}}}}\Phi^{-1}(\epsilon_{\mathrm{e}})$ is an achievable $(0,\epsilon_{\mathrm{e}})_{\mathrm{ave,ave}}$-erasure second-order coding rate for W^n. Indeed, here the total error $\epsilon_{\mathrm{t}}=\epsilon_{\mathrm{e}}$. However, any achievable $(0,\epsilon_{\mathrm{t}})_{\mathrm{ave,ave}}$-second-order coding rate is also an achievable $(\epsilon_{\mathrm{u}},\epsilon_{\mathrm{t}})_{\mathrm{ave,ave}}$-second-order coding rate for any $\epsilon_{\mathrm{u}}\in[0,\epsilon_{\mathrm{t}})$. Furthermore, by an expurgation argument similar to (55), the same statement can be proved under the $(\max,\max)$ setting.

For this proof, we show that the second-order capacity in (23) is also *universally attainable*–i.e. the code does not require channel knowledge. A simpler proof based on thresholding the likelihood can also be be used; however, channel knowledge is required. See Sec. V-B3 for an analogue of the alternative strategy.

Fix a type $P\in\mathcal{P}_n(\mathcal{X})$. Generate M_n codewords uniformly at random from the type class $\mathcal{T}_P$. Denote the random codebook as $\{\mathbf{X}(m): m\in[M_n]\}\subset\mathcal{T}_P$. The number M_n is to be chosen later. Let $\gamma\geq 0$ be some threshold to be chosen later. At the receiver, we decode to $\hat{m}$ if and only if $\hat{m}$ is the unique message to satisfy

$$\hat{I}(\mathbf{X}(\hat{m})\wedge\mathbf{Y}) \geq \gamma \tag{56}$$

where $\hat{I}(\mathbf{x}\wedge\mathbf{y})$ is the empirical mutual information of $(\mathbf{x},\mathbf{y})$, i.e. the mutual information of the random variables $(\tilde{X},\tilde{Y})$ whose distribution is the joint type $P_{\mathbf{x},\mathbf{y}}$. Assume as usual that the true message $m=1$. We use the following elementary result which is shown in the proof of the packing lemma [19, Lem. 10.1].

Lemma 5. *Let $P\in\mathcal{P}_n(\mathcal{X})$ be any n-type. Let $\mathbf{X}$ and $\overline{\mathbf{X}}$ be selected independently and uniformly at random from the type class $\mathcal{T}_P\subset\mathcal{X}^n$. Let $\mathbf{Y}$ be the channel output when $\mathbf{X}$ is the input, i.e. $\mathbf{Y}|\{\mathbf{X}=\mathbf{x}\}\sim\prod_{j=1}^n W(\cdot|x_j)$. Then, for every $\gamma>0$ and every $n\in\mathbb{N}$,*

$$\Pr\left[\hat{I}(\overline{\mathbf{X}}\wedge\mathbf{Y})>\gamma\right] \leq (n+1)^{|\mathcal{X}|+|\mathcal{X}||\mathcal{Y}|}\exp(-n\gamma). \tag{57}$$

The undetected error probability is bounded as

$$\begin{aligned}
\Pr[\mathcal{E}_{\mathrm{u}}] &\leq \Pr\left[\max_{m\geq 2}\hat{I}(\mathbf{X}(m)\wedge\mathbf{Y})\geq\gamma\right] && (58)\\
&\leq (M_n-1)\Pr\left[\hat{I}(\mathbf{X}(2)\wedge\mathbf{Y})\geq\gamma\right] && (59)\\
&\leq (M_n-1)(n+1)^{|\mathcal{X}||\mathcal{Y}|+|\mathcal{X}|}\exp(-n\gamma), && (60)
\end{aligned}$$

where (59) follows from the union bound and the fact that the codewords are generated in an identical manner, and (60) follows from Lemma 5 noting that $\mathbf{X}(2)$ is independent of $\mathbf{Y}$, the channel output when $\mathbf{X}(1)$ is the input.

We let the random conditional type of $\mathbf{Y}$ given $\mathbf{X}(1)$ be U. The erasure probability can be bounded as

$$\Pr[\mathcal{E}_{\mathrm{e}}] = \Pr\left[\hat{I}(\mathbf{X}(1) \wedge \mathbf{Y}) \le \gamma\right] \tag{61}$$

$$= \Pr\left[I(P,W) + \sum_{x,y}(U(y|x) - W(y|x))I'_W(y|x) + O(\|U-W\|^2) \le \gamma\right] \tag{62}$$

where the final step follows by Taylor expanding $U \mapsto I(P,U)$ around $U = W$ and $I'_W(y|x) := \frac{\partial I(P,U)}{\partial U(y|x)}\big|_{U=W}$. We also can bound the remainder term uniformly [31] yielding

$$\Pr[\mathcal{E}_{\mathrm{e}}] \le \Pr\left[I(P,W) + \sum_{x,y}(U(y|x) - W(y|x))I'_W(y|x) \le \gamma + O\left(\frac{\log n}{n}\right)\right] + O(n^{-2}). \tag{63}$$

Wang-Ingber-Kochman [31] computed the relevant first-, second- and third-order statistics of the random variable $\sum_{x,y}(U(y|x)-W(y|x))I'_W(y|x)$ allowing us to apply the Berry-Esseen theorem [32, Ch. XVI.5] to the probability in (63), leading to

$$\Pr[\mathcal{E}_{\mathrm{e}}] \le \Phi\left(\frac{\gamma + O(\frac{\log n}{n}) - I(P,W)}{\sqrt{V(P,W)/n}}\right) + O(n^{-1/2}). \tag{64}$$

Note that the implied constants in the $O(\,\cdot\,)$-notation in (64) are bounded because of the discreteness of the alphabets. Hence, we set

$$\gamma = I(P,W) + \sqrt{\frac{V(P,W)}{n}}\Phi^{-1}(\epsilon'_{\mathrm{e}}) \tag{65}$$

to assert that

$$\Pr[\mathcal{E}_{\mathrm{e}}] \le \epsilon'_{\mathrm{e}} + O(n^{-1/2}). \tag{66}$$

We then set M_n to be the smallest integer satisfying

$$\log M_n \ge n\gamma - \left(|\mathcal{X}||\mathcal{Y}| + |\mathcal{X}| + \frac{1}{2}\right)\log n. \tag{67}$$

Then we may assert from (60) that

$$\Pr[\mathcal{E}_{\mathrm{u}}] \le n^{-1/2}. \tag{68}$$

Hence, we have that

$$\log M_n \ge nI(P,W) + \sqrt{nV(P,W)}\Phi^{-1}(\epsilon'_{\mathrm{e}}) + O\left(\log n\right). \tag{69}$$

Let P_X^* achieve $V_{\epsilon_{\mathrm{e}}}$. This means that $P_X^* \in \arg\min_{P\in\Pi} V(P,W)$ if $\epsilon_{\mathrm{e}} < 1/2$ and $P_X^* \in \arg\max_{P\in\Pi} V(P,W)$ if $\epsilon_{\mathrm{e}} \ge 1/2$. By choosing P to be an n-type that is the closest to P_X^* (i.e., $P \in \arg\min_{P\in\mathcal{P}_n(\mathcal{X})}\|P - P_X^*\|_1$), we obtain, by the usual approximation arguments [13, Lem. 7],

$$\log M_n \ge nC + \sqrt{nV_{\epsilon'_{\mathrm{e}}}}\Phi^{-1}(\epsilon'_{\mathrm{e}}) + O\left(\log n\right). \tag{70}$$

We have proved the the random ensemble satisfies (66) and (68) but it is not clear yet there exists a single deterministic code that satisfies the same two bounds. To show this, let $\theta \in (0,1)$. Set $\epsilon_{\mathrm{u}} := \frac{1}{\theta}n^{-1/2}$ and $\epsilon_{\mathrm{e}} := \frac{1}{1-\theta}(\epsilon'_{\mathrm{e}} + O(n^{-1/2}))$ where $\epsilon'_{\mathrm{e}} + O(n^{-1/2})$ denotes the right-hand-side of (66). Then, making the expectation over the random code $\mathcal{C}$ explicit, (66) and (68) can be written as

$$\mathrm{E}[\Pr[\mathcal{E}_{\mathrm{e}}|\mathcal{C}]] \le (1-\theta)\epsilon_{\mathrm{e}}, \quad \text{and} \quad \mathrm{E}[\Pr[\mathcal{E}_{\mathrm{u}}|\mathcal{C}]] \le \theta\epsilon_{\mathrm{u}}. \tag{71}$$

Put $\eta := \theta/2$. By Markov's inequality,

$$\Pr[\mathcal{A}] \le 1-\theta, \quad \text{and} \quad \Pr[\mathcal{B}] \le \theta - \eta, \tag{72}$$

where event $\mathcal{A} := \{\Pr[\mathcal{E}_{\mathrm{e}}|\mathcal{C}] > \frac{1}{1-\theta}\mathrm{E}[\Pr[\mathcal{E}_{\mathrm{e}}|\mathcal{C}]]\}$ and event $\mathcal{B} := \{\Pr[\mathcal{E}_{\mathrm{u}}|\mathcal{C}] > \frac{1}{\theta-\eta}\mathrm{E}[\Pr[\mathcal{E}_{\mathrm{u}}|\mathcal{C}]]\}$. This implies that there exists a *deterministic* code $\mathscr{C}_n^*$ for which

$$\Pr[\mathcal{E}_{\mathrm{e}}|\mathscr{C}_n^*] \le \frac{1}{1-\theta}\mathrm{E}[\Pr[\mathcal{E}_{\mathrm{e}}|\mathcal{C}]] \le \epsilon_{\mathrm{e}}, \quad \text{and} \tag{73}$$

$$\Pr[\mathcal{E}_{\mathrm{u}}|\mathscr{C}_n^*] \le \frac{1}{\theta-\eta}\mathrm{E}[\Pr[\mathcal{E}_{\mathrm{u}}|\mathcal{C}]] \le \frac{\theta}{\theta-\eta}\epsilon_{\mathrm{u}} = \frac{2}{\theta\sqrt{n}}. \tag{74}$$

Letting $n \to \infty$, we conclude there exists a sequence of deterministic codes $\{\mathscr{C}_n^*\}_{n\geq 1}$ such that

$$\limsup_{n\to\infty} \Pr[\mathcal{E}_{\mathrm{e}}|\mathscr{C}_n^*] \leq \epsilon_{\mathrm{e}}, \quad \text{and} \quad \lim_{n\to\infty} \Pr[\mathcal{E}_{\mathrm{u}}|\mathscr{C}_n^*] = 0. \tag{75}$$

Now, according to (70) and the relation between ϵ_{e} and ϵ'_{e},

$$\liminf_{n\to\infty} \frac{1}{\sqrt{n}}(\log M_n - nC) \geq \sqrt{V_{(1-\theta)\epsilon_{\mathrm{e}}}}\Phi^{-1}((1-\theta)\epsilon_{\mathrm{e}}). \tag{76}$$

Finally take $\theta \downarrow 0$ to complete the proof.

B. List Decoding: Proof of Theorem 2

We first prove the non-asymptotic converse bound for list decoding (Proposition 3) in Sec. V-B1. Subsequently, we prove (35)–(37) in unison in both the converse (Sec. V-B2) and direct (Sec. V-B3) parts.

1) Proof of Proposition 3: We modify the argument leading to [13, Prop. 6], which was inspired by [33], [34], so that it is applicable to the list decoding setting. Let P and Q be two probability measures on the same space $\mathcal{Z}$. Let $\delta : \mathcal{Z} \to \{0,1\}$ represent a (deterministic) hypothesis test between P and Q where $\delta = 1$ implies deciding in favor of P. We define the *ϵ-hypothesis testing divergence*

$$D_{\mathrm{h}}^{\epsilon}(P\|Q) := \sup\Big\{R \in \mathbb{R} : \exists\, \delta : \mathcal{Z} \to \{0,1\}, \mathrm{E}_Q[\delta(Z)] \leq (1-\epsilon)\exp(-R), \mathrm{E}_P[\delta(Z)] \geq 1-\epsilon\Big\} \tag{77}$$

Note that the ϵ-hypothesis testing divergence is related to $\beta_{1-\epsilon}$ as follows:

$$D_{\mathrm{h}}^{\epsilon}(P\|Q) = \log(1-\epsilon) - \log \beta_{1-\epsilon}(P,Q). \tag{78}$$

Some properties of $D_{\mathrm{h}}^{\epsilon}$ include (i) $D_{\mathrm{h}}^{\epsilon}(P\|Q) \geq 0$ with equality iff $P = Q$ almost everywhere; (ii) For any random transformation $U : \mathcal{Z} \to \mathcal{Z}'$, the data processing inequality [13, Prop. 6] implies that $D_{\mathrm{h}}^{\epsilon}(PU\|QU) \leq D_{\mathrm{h}}^{\epsilon}(P\|Q)$.

Any $(|\mathcal{M}|, L, \epsilon)_{\mathrm{ave}}$-list code (f, φ) with message set $\mathcal{M}$ (Definition 5) induces the Markov chain $M - X - Y - \mathcal{S}_L$ where M is uniform on $\mathcal{M}$ and $\mathcal{S}_L \in \binom{\mathcal{M}}{\leq L}$ is a random subset of $\mathcal{M}$ with size not exceeding L. This defines the joint distribution $P_{MXY\mathcal{S}_L}(m,x,y,s_L) = \frac{1}{|\mathcal{M}|}\mathbf{1}\{x = f(m)\}W(y|x)\mathbf{1}\{\varphi(m) \in s_L\}$. Fix $Q \in \mathcal{P}(\mathcal{Y})$. Due to the data processing inequality for $D_{\mathrm{h}}^{\epsilon}$, we have

$$D_{\mathrm{h}}^{\epsilon}(P \times W\|P \times Q) = D_{\mathrm{h}}^{\epsilon}(P_{XY}\|P_X \times Q_Y) \geq D_{\mathrm{h}}^{\epsilon}(P_{M\mathcal{S}_L}\|P_M \times Q_{\mathcal{S}_L}) \tag{79}$$

where $P_X = P$ and $Q_{\mathcal{S}_L}$ is the distribution on the subsets of $\mathcal{M}$ of size not exceeding L induced by the decoder applied to $Q_Y = Q$. Moreover, consider the test (with the identification of the space $\mathcal{Z} \leftarrow \mathcal{M} \times \binom{\mathcal{M}}{\leq L}$ in (77))

$$\delta(M, \mathcal{S}_L) := \mathbf{1}\{M \in \mathcal{S}_L\}. \tag{80}$$

Then under the null hypothesis $P_{M\mathcal{S}_L}$, we have

$$\mathrm{E}_{P_{M\mathcal{S}_L}}[\delta(M,\mathcal{S}_L)] = \Pr[M \in \mathcal{S}_L] \geq 1-\epsilon \tag{81}$$

because the average error probability of the list code is no larger than ϵ. In addition, under the alternate hypothesis $P_M \times Q_{\mathcal{S}_L}$,

$$\begin{aligned}
\mathrm{E}_{P_M \times Q_{\mathcal{S}_L}}[\delta(M,\mathcal{S}_L)] &= \sum_{m\in\mathcal{M}} \frac{1}{|\mathcal{M}|} \sum_{\mathscr{S}\in\binom{\mathcal{M}}{\leq L}} Q_{\mathcal{S}_L}(\mathscr{S})\delta(m,\mathscr{S}) && (82)\\
&= \frac{1}{|\mathcal{M}|} \sum_{\mathscr{S}\in\binom{\mathcal{M}}{\leq L}} Q_{\mathcal{S}_L}(\mathscr{S}) \sum_{m\in\mathcal{M}} \delta(m,\mathscr{S}) && (83)\\
&= \frac{1}{|\mathcal{M}|} \sum_{\mathscr{S}\in\binom{\mathcal{M}}{\leq L}} Q_{\mathcal{S}_L}(\mathscr{S})|\mathscr{S}| && (84)\\
&\leq \frac{L}{|\mathcal{M}|} \sum_{\mathscr{S}\in\binom{\mathcal{M}}{\leq L}} Q_{\mathcal{S}_L}(\mathscr{S}) && (85)\\
&= \frac{L}{|\mathcal{M}|}, && (86)
\end{aligned}$$

where (85) follows because $|\mathscr{S}| \leq L$ for all subsets $\mathscr{S} \in \binom{\mathcal{M}}{\leq L}$. Hence, by the definition of the ϵ-hypothesis testing divergence $D_{\mathrm{h}}^{\epsilon}$ in (77), we have

$$D_{\mathrm{h}}^{\epsilon}(P_{M\mathcal{S}_L} \| P_M \times Q_{\mathcal{S}_L}) \geq \log \frac{|\mathcal{M}|}{L} + \log(1-\epsilon). \tag{87}$$

By uniting (79) and (87), maximizing over $P \in \mathcal{P}(\mathcal{X})$ to make the bound code independent, and minimizing over $Q \in \mathcal{P}(\mathcal{Y})$ which was arbitrary, we obtain

$$\log \frac{|\mathcal{M}|}{L} \leq \inf_{Q \in \mathcal{P}(\mathcal{Y})} \sup_{P \in \mathcal{P}(\mathcal{X})} D_{\mathrm{h}}^{\epsilon}(P \times W \| P \times Q) + \log \frac{1}{1-\epsilon} \tag{88}$$

which is exactly the same as (38) in Proposition 3.

2) Converse Part: We assume that $V_\epsilon > 0$. Now, we may choose $Q_{\mathbf{Y}} \in \mathcal{P}(\mathcal{Y}^n)$ as in [13, Sec III.C] and follow the analysis in [13, Props. 6, 8 and 10(i)] to upper bound $D_{\mathrm{h}}^{\epsilon}(P_{\mathbf{X}} \times W^n \| P_{\mathbf{X}} \times Q_{\mathbf{Y}})$ by relaxing it to a quantity known as the ϵ-information spectrum divergence. This yields

$$\log \frac{M_n}{L_n} \leq nC + \sqrt{nV_\epsilon}\Phi^{-1}(\epsilon) + \frac{1}{2}\log n + O(1). \tag{89}$$

To prove (35), notice that $\log L_n = O(\sqrt{n}\, l)$ so the second-order term $r^*_{\mathrm{list,a}}(l,\epsilon) \leq l + \sqrt{nV_\epsilon}\Phi^{-1}(\epsilon)$. For (36), $L_n = O(n^\alpha)$, so the third-order term $s^*_{\mathrm{list}}(\alpha,\epsilon) \leq \frac{1}{2} + \alpha$ as desired.

To prove the stronger statement concerning symmetric and singular DMCs in (37), we note that per [30, Thm. 22], the saddle-point in (88) is attained for $P_{\mathbf{X}}$ being the uniform distribution on $\mathcal{X}^n$. We also use the output distribution suggested by Altuğ and Wagner in [14, Sec. IV.A]. This shows that $D_{\mathrm{h}}^{\epsilon}(P_{\mathbf{X}} \times W^n \| P_{\mathbf{X}} \times Q_{\mathbf{Y}})$ is bounded above by the Gaussian approximation plus a constant term, completing the proof of $s^*_{\mathrm{list}}(\alpha,\epsilon) \leq \alpha$ in view of the fact that $L_n = O(n^\alpha)$.

3) Direct Part: To get the third-order result for the lower bound in (36), we use i.i.d. random codes with distribution P_X^* achieving V_ϵ. We also use threshold decoding of the information density. Contrast this to using constant composition codes and maximum empirical mutual information decoding in Sec. V-A2 which generally results in worse third-order terms. Now the decoder outputs *all* messages m whose log-likelihood ratio exceeds $\gamma \geq 0$, i.e. the list is

$$\mathcal{L} := \left\{ m \in [M_n] : \frac{1}{n} \log \frac{W^n(\mathbf{Y}|\mathbf{X}(m))}{(P_X^* W)^n(\mathbf{Y})} \geq \gamma \right\}. \tag{90}$$

We now analyze the error probability assuming that message $m = 1$ was sent. The two error events are

$$\mathcal{E}_1 := \left\{ \frac{1}{n} \log \frac{W^n(\mathbf{Y}|\mathbf{X}(1))}{(P_X^* W)^n(\mathbf{Y})} < \gamma \right\}, \quad \text{and} \tag{91}$$

$$\mathcal{E}_2 := \{ |\mathcal{L}| > L_n \}. \tag{92}$$

The probability of $\mathcal{E}_1$ can be analyzed using the Berry-Esseen theorem [32, Ch. XVI.5]. This yields

$$\Pr[\mathcal{E}_1] \leq \Phi\left(\frac{\gamma - C}{\sqrt{V_\epsilon/n}} \right) + O(n^{-1/2}). \tag{93}$$

The ϵ-dispersion V_ϵ appears in the denominator in (93) because the unconditional information variance equals the conditional information variance for capacity-achieving distributions [10, Lem. 62] and P_X^* is chosen to achieve V_ϵ. Now, consider the expectation of the size of the list of incorrect messages $\mathcal{L}_{\setminus 1} := \mathcal{L} \setminus \{1\}$. Indeed

$$\mathrm{E}\big[|\mathcal{L}_{\setminus 1}|\big] = \mathrm{E}\left[\sum_{m=2}^{M_n} \mathbf{1}\left\{ \frac{1}{n} \log \frac{W^n(\mathbf{Y}|\mathbf{X}(m))}{(P_X^* W)^n(\mathbf{Y})} \geq \gamma \right\} \right] \tag{94}$$

$$= \sum_{m=2}^{M_n} \Pr\left[\frac{1}{n} \log \frac{W^n(\mathbf{Y}|\mathbf{X}(m))}{(P_X^* W)^n(\mathbf{Y})} \geq \gamma \right]. \tag{95}$$

Now, we analyze the probability in (95). By symmetry, all summands in (95) are identical so we just focus on the $m = 2$ term. By [12, Prop. 6.2], which follows from a strong large deviations result of Chaganty and Sethuraman [35, Thms. 3.3 and 3.5], we have that[3]

$$\Pr\left[\frac{1}{n}\log\frac{W^n(\mathbf{Y}|\mathbf{X}(2))}{(P_X^* W)^n(\mathbf{Y})} \geq \gamma\right] \sim \frac{a\,\exp(-n\gamma)}{\sqrt{n}} \tag{96}$$

for some channel dependent constant $a > 0$. A non-asymptotic upper bound can also be derived using [10, Lem. 47]. Note that γ can be a sequence in n under some conditions [12, Prop. 6.2]. It can easily be verified that our choice of γ in (100) satisfies such conditions. Thus plugging (96) into (95) and bounding $M_n - 1$ by M_n we have that

$$\mathrm{E}\big[|\mathcal{L}_{\setminus 1}|\big] \leq \frac{M_n b\exp(-n\gamma)}{\sqrt{n}} \tag{97}$$

for some channel dependent constant $b > 0$ and n sufficiently large. Since $|\mathcal{L}| \leq |\mathcal{L}_{\setminus 1}| + 1$ (consider the two different cases in which $1 \in \mathcal{L}$ and otherwise), the probability of $\mathcal{E}_2$ can be bounded as

$$\Pr[\mathcal{E}_2] = \Pr\left[|\mathcal{L}| > L_n\right] \leq \Pr\left[|\mathcal{L}_{\setminus 1}| > L_n - 1\right]. \tag{98}$$

By Markov's inequality,

$$\Pr[\mathcal{E}_2] \leq \frac{\mathrm{E}\big[|\mathcal{L}_{\setminus 1}|\big]}{L_n - 1} \leq \frac{M_n b\exp(-n\gamma)}{\sqrt{n}\,(L_n - 1)}. \tag{99}$$

Now choose

$$\gamma = C + \sqrt{\frac{V_\epsilon}{n}}\Phi^{-1}(\epsilon) \tag{100}$$

so $\Pr[\mathcal{E}_1] \leq \epsilon + O(n^{-1/2})$ according to (93). Also choose

$$M_n = \lfloor (L_n - 1)\exp(n\gamma)\rfloor \tag{101}$$

so $\Pr[\mathcal{E}_2|\mathcal{E}_1^c] \leq O(n^{-1/2})$ according to (99). As such the total error probability $\Pr[\mathcal{E}_1 \cup \mathcal{E}_2]$ (probability that the true message does not belong to the list or the list size exceeds L_n) is bounded from above by $\epsilon + O(n^{-1/2})$, which satisfies (the more stringent condition in) (21).

In the second-order case in which $\log L_n = O(\sqrt{n}\,l)$, the choice of M_n in (101) yields $r^*_{\mathrm{list,a}}(l,\epsilon) \geq l + \sqrt{nV_\epsilon}\Phi^{-1}(\epsilon)$, proving the direct part of (35). In the third-order case in which $L_n = O(n^\alpha)$, the choice of M_n in (101) yields $s^*_{\mathrm{list}}(\alpha,\epsilon) \geq \alpha$, proving the direct part of (36).

Note that unlike the erasures setting, in this case we do not need to augment the proof with the argument involving Markov's inequality (cf. argument leading to (74)) because here, there is only a single error criterion.

C. Slepian-Wolf with Erasure Option: Proof of Theorem 4

1) Converse part: The technique developed by Miyake and Kanaya [28, Sec. 4.2] applies directly. This allows us to conclude that every $(M_n, \epsilon_{\mathrm{u},n}, \epsilon_{\mathrm{t},n})$-code for the correlated source must satisfy

$$\epsilon_{\mathrm{t},n} \geq \Pr\left(\frac{1}{n}\log\frac{1}{P_{\mathbf{X}|\mathbf{Y}}(\mathbf{X}|\mathbf{Y})} \geq \frac{1}{n}\log M_n + \gamma\right) + \exp(-n\gamma), \tag{102}$$

for any $\gamma > 0$ regardless of $\epsilon_{\mathrm{u},n} \in [0, \epsilon_{\mathrm{t},n})$. We immediately obtain the converse by setting $\gamma = \frac{\log n}{2n}$ and applying the Berry-Esseen theorem [32, Ch. XVI.5] to the probability in (102). See the proof of the converse of [16, Thm. 1] for details.

[3]The notation $a_n \sim b_n$ means that the limit of the ratio of a_n and b_n tends to one as n tends to infinity.

2) Direct part: By the same argument as the channel coding case in Sec. V-A2, it suffices to show that $\sqrt{V}\Phi^{-1}(1-\epsilon_{\mathrm{t}})$ is $(0,\epsilon_{\mathrm{t}})$-achievable.

As in Sec. V-A2, we provide a universal coding scheme. Randomly and independently partition the set of source sequences $\mathcal{X}^n$ into $M_n = \lceil \exp(nR_n) \rceil$ bins where the "rate" R_n is to be chosen later. Let $\mathcal{B}(m)$ be the set of source sequences $\mathbf{x}$ which are (randomly) mapped to the bin indexed by $m \in [M_n]$. If $\mathbf{x}$ is received by the encoder, send message m.

Let $\hat{H}(\mathbf{x}|\mathbf{y})$ be the empirical conditional entropy of the sequences $(\mathbf{x},\mathbf{y})$, i.e., $\hat{H}(\mathbf{x}|\mathbf{y}) = H(U|P_{\mathbf{y}})$ where U is the conditional type of $\mathbf{x}$ given $\mathbf{y}$. The decoder, given $\mathbf{y} \in \mathcal{Y}^n$ and $m \in [M_n]$, finds a source sequence $\hat{\mathbf{x}}$ in $\mathcal{B}(m)$ satisfying

$$\hat{H}(\hat{\mathbf{x}}|\mathbf{y}) \le \gamma \tag{103}$$

for some threshold γ, which will be chosen later. If there is no such source sequence in bin $\mathcal{B}(m)$, declare an erasure event. If there is more than one such source sequence, uniformly pick one.

Let $\mathbf{X}$ and $\mathbf{Y}$ be the randomly generated sequences from $\prod_{j=1}^n P_{XY}(x_j,y_j)$. Because an erasure event implies that $\hat{H}(\mathbf{X}|\mathbf{Y}) > \gamma$, the probability of erasure can be bounded as

$$\begin{aligned}
\Pr[\mathcal{E}_{\mathrm{e}}] &\le \Pr\left(\hat{H}(\mathbf{X}|\mathbf{Y}) > \gamma\right) && (104)\\
&\le \Phi\left(\frac{H(X|Y)-\gamma}{\sqrt{V(X|Y)/n}}\right) + O(n^{-1/2}), && (105)
\end{aligned}$$

where the last step follows by Taylor expanding $\hat{H}(\mathbf{X}|\mathbf{Y})$ around $H(X|Y)$ and then applying the Berry-Esseen theorem [32, Ch. XVI.5]. See the so-called vector rate redundancy theorem in [16] for details. The probability of undetected error can be written as

$$\Pr[\mathcal{E}_{\mathrm{u}}] = \Pr\left[\exists\, \tilde{\mathbf{x}} \in \mathcal{B}(M) \setminus \{\mathbf{X}\} : \hat{H}(\tilde{\mathbf{x}}|\mathbf{Y}) \le \gamma\right], \tag{106}$$

where M is the bin index corresponding to the random source $\mathbf{X}$. By symmetry, it is enough to fix the bin index to be $M=1$. Now, we condition on various values of $(\mathbf{x},\mathbf{y}) \in \mathcal{X}^n \times \mathcal{Y}^n$ as follows:

$$\begin{aligned}
\Pr[\mathcal{E}_{\mathrm{u}}] &= \sum_{\mathbf{x},\mathbf{y}} P_{XY}^n(\mathbf{x},\mathbf{y}) \Pr\left[\exists\, \tilde{\mathbf{x}} \in \mathcal{B}(1) \setminus \{\mathbf{x}\} : \hat{H}(\tilde{\mathbf{x}}|\mathbf{Y}) \le \gamma \,\middle|\, \mathbf{X}=\mathbf{x}, \mathbf{Y}=\mathbf{y}\right] && (107)\\
&\le \sum_{\mathbf{x},\mathbf{y}} P_{XY}^n(\mathbf{x},\mathbf{y}) \sum_{\tilde{\mathbf{x}} \neq \mathbf{x}: \hat{H}(\tilde{\mathbf{x}}|\mathbf{y}) \le \gamma} \Pr\left[\tilde{\mathbf{x}} \in \mathcal{B}(1)\right] && (108)\\
&= \sum_{\mathbf{x},\mathbf{y}} P_{XY}^n(\mathbf{x},\mathbf{y}) \sum_{U \in \mathcal{U}_n(\mathcal{X};P_{\mathbf{y}}): H(U|P_{\mathbf{y}}) \le \gamma} \sum_{\tilde{\mathbf{x}} \in \mathcal{T}_U(\mathbf{y})} \Pr\left[\tilde{\mathbf{x}} \in \mathcal{B}(1)\right] && (109)\\
&\le \sum_{\mathbf{x},\mathbf{y}} P_{XY}^n(\mathbf{x},\mathbf{y}) \sum_{U \in \mathcal{U}_n(\mathcal{X};P_{\mathbf{y}}): H(U|P_{\mathbf{y}}) \le \gamma} |\mathcal{T}_U(\mathbf{y})| \exp(-nR_n) && (110)\\
&\le \sum_{\mathbf{x},\mathbf{y}} P_{XY}^n(\mathbf{x},\mathbf{y}) \sum_{U \in \mathcal{U}_n(\mathcal{X};P_{\mathbf{y}}): H(U|P_{\mathbf{y}}) \le \gamma} \exp(nH(U|P_{\mathbf{y}})) \exp(-nR_n) && (111)\\
&\le \sum_{\mathbf{x},\mathbf{y}} P_{XY}^n(\mathbf{x},\mathbf{y}) \sum_{U \in \mathcal{U}_n(\mathcal{X};P_{\mathbf{y}}): H(U|P_{\mathbf{y}}) \le \gamma} \exp(n\gamma) \exp(-nR_n) && (112)\\
&\le (n+1)^{|\mathcal{X}||\mathcal{Y}|} \exp(-n(R_n-\gamma)), && (113)
\end{aligned}$$

where (108) follows from the fact that the binning is independent of the generation of $\mathbf{x},\mathbf{y}$, in (109) we partitioned $\tilde{\mathbf{x}}$ into U-shells compatible with the type of $\mathbf{y}$, (110) follows from the uniformity in binning and the fact that the number of bins is $M_n = \lceil \exp(nR_n) \rceil$, (111) follows from the fact that $|\mathcal{T}_U(\mathbf{y})| \le \exp(nH(U|P_{\mathbf{y}}))$, (112) uses the fact that $H(U|P_{\mathbf{y}}) \le \gamma$, and finally, (113) follows from the type counting lemma [19, Lem. 2.1].

We now choose

$$\gamma = H(X|Y) + \sqrt{\frac{V(X|Y)}{n}}\Phi^{-1}(1-\epsilon_{\mathrm{t}}) \tag{114}$$

so (105) is $\epsilon_{\mathrm{t}} + O(n^{-1/2})$. We also pick

$$R_n = \gamma + (|\mathcal{X}||\mathcal{Y}| + 1/2)\frac{\log(n+1)}{n} \tag{115}$$

so (113) is of the order $O(n^{-1/2})$. This, together with the Markov inequality argument at the end of the proof in Section V-A2 completes the proof of the direct part because $\sqrt{V(X|Y)}\Phi^{-1}(1-\epsilon_{\mathrm{t}})$ is a $(0, \epsilon_{\mathrm{t}})$-achievable second-order coding rate for the Slepian-Wolf problem with correlated source P_{XY}.

Acknowledgements: VT would like to thank Yanina Shkel for sharing her MATLAB code to generate the finite blocklength performance bounds in Fig. 1. He would also like to acknowledge helpful discussions with Masahito Hayashi, Marco Tomamichel and Jonathan Scarlett.

REFERENCES

[1] G. D. Forney. Exponential error bounds for erasure, list, and decision feedback schemes. *IEEE Trans. on Inf. Th.*, 14:206–220, 1968.
[2] C. E. Shannon, R. G. Gallager, and E. R. Berlekamp. Lower bounds to error probability for coding in discrete memoryless channels I-II. *Information and Control*, 10:65–103,522–552, 1967.
[3] I. E. Telatar and R. G. Gallager. New exponential upper bounds to error and erasure probabilities. In *Intl. Symp. on Inf. Th.*, page 379, Trondheim, Norway, 1994.
[4] V. M. Blinovsky. Error probability exponent of list decoding at low rates. *Problems of Information Transmission*, 27(4):277–287, 2001.
[5] P. Moulin. A Neyman-Pearson approach to universal erasure and list decoding. *IEEE Trans. on Inf. Th.*, 55:4462–4478, Oct 2009.
[6] N. Merhav. Error exponents of erasure/list decoding revisited via moments of distance enumerators. *IEEE Trans. on Inf. Th.*, 54:4439–4447, Oct 2008.
[7] N. Merhav. List decoding–random coding exponents and expurgated exponents. *submitted to the IEEE Trans. on Inf. Th.*, 2013. arXiv:1311.7298.
[8] V. Strassen. Asymptotische Abschätzungen in Shannons Informationstheorie. In *Trans. Third Prague Conf. Inf. Theory*, pages 689–723, Prague, 1962. Available at http://www.math.cornell.edu/∼pmlut/strassen.pdf.
[9] J. H. B. Kemperman. Studies in Coding Theory I. Technical report, University of Rochester, NY, 1962. Available at https://www.dropbox.com/s/amrio27pea0vz1f/Kemperman.pdf.
[10] Y. Polyanskiy, H. V. Poor, and S. Verdú. Channel coding rate in the finite blocklength regime. *IEEE Trans. on Inf. Th.*, 56:2307–2359, May 2010.
[11] M. Hayashi. Information spectrum approach to second-order coding rate in channel coding. *IEEE Trans. on Inf. Th.*, 55:4947–4966, Nov 2009.
[12] P. Moulin. The log-volume of optimal codes for memoryless channels, within a few nats. Nov 2013. arXiv:1311.0181.
[13] M. Tomamichel and V. Y. F. Tan. A tight upper bound for the third-order asymptotics of most discrete memoryless channels. *IEEE Trans. on Inf. Th.*, 59(11):7041–7051, Nov 2013.
[14] Y. Altuğ and A. Wagner. The third-order term in the normal approximation for singular channels. 2013. arXiv:309.5126.
[15] V. Y. F. Tan and M. Tomamichel. The third-order term in the normal approximation for the AWGN channel. 2013. arXiv:1311.2237v2.
[16] V. Y. F. Tan and O. Kosut. On the dispersions of three network information theory problems. *IEEE Trans. on Inf. Th.*, 60(2):881–903, Feb 2014.
[17] D. Slepian and J. K. Wolf. Noiseless coding of correlated information sources. *IEEE Trans. on Inf. Th.*, 19:471–80, 1973.
[18] R. G. Gallager. *Information Theory and Reliable Communication*. Wiley, New York, 1968.
[19] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Cambridge University Press, 2011.
[20] N. Merhav. Erasure/list exponents for Slepian-Wolf decoding. *submitted to the IEEE Trans. on Inf. Th.*, 2013. arXiv:1305.5626.
[21] Y. Polyanskiy, H. V. Poor, and S. Verdú. Feedback in the non-asymptotic regime. *IEEE Trans. on Inf. Th.*, 57(8):4903–4925, 2011.
[22] T.-Y. Chen, A. R. Williamson, N. Seshadri, and R. D. Wesel. Feedback communication systems with limitations on incremental redundancy. *submitted to the IEEE Trans. on Inf. Th.*, 2013. arXiv:1309.0707.
[23] A. R. Williamson, T.-Y. Chen, and R. D. Wesel. Reliability-based error detection for feedback communication with low latency. In *Intl. Symp. Inf. Th.*, Istanbul, Turkey, 2013.
[24] Y. Polyanskiy. *Channel coding: Non-asymptotic fundamental limits*. PhD thesis, Princeton University, 2010.
[25] R. Nomura and T. S. Han. Second-order Slepian-Wolf coding theorems for non-mixed and mixed sources. In *Int. Symp. Inf. Th.*, Istanbul, Turkey, 2013. `arXiv:1207.2505 [cs.IT]`.
[26] S. Verdú and I. Kontoyiannis. Optimal lossless data compression: Non-asymptotics and asymptotics. *IEEE Trans. on Inf. Th.*, 60(2):777–795, Feb 2014.
[27] T. M. Cover. A proof of the data compression theorem of Slepian and Wolf for ergodic sources. *IEEE Trans. Inf. Th.*, 21(3):226–228, March 1975.
[28] S. Miyake and F. Kanaya. Coding theorems on correlated general sources. *IEICE Trans. on Fundamentals of Electronics, Communications and Computer*, E78-A(9):1063–70, 1995.
[29] T. S. Han. *Information-Spectrum Methods in Information Theory*. Springer Berlin Heidelberg, Feb 2003.
[30] Y. Polyanskiy. Saddle point in the minimax converse for channel coding. *IEEE Trans. on Inf. Th.*, 59:2576–2595, May 2013.
[31] D. Wang, A. Ingber, and Y. Kochman. The dispersion of joint source-channel coding. In *Allerton Conference*, 2011. arXiv:1109.6310.
[32] W. Feller. *An Introduction to Probability Theory and Its Applications*. John Wiley and Sons, 2nd edition, 1971.
[33] L. Wang, R. Colbeck, and R. Renner. Simple channel coding bounds. In *Intl. Symp. Inf. Th.*, Seoul, South Korea, 2009.
[34] L. Wang and R. Renner. One-shot classical-quantum capacity and hypothesis testing. *Physical Review Letters*, 108:200501, May 2012.
[35] N. R. Chaganty and J. Sethuraman. Strong large deviation and local limit theorems. *Annals of Probability*, 21(3):1671–1690, 1993.